

08-15-07

Serial No. 10/620,080
Page 1 of 17

AF
JRW

Attorney Docket No.: ASSIA 20.502 (056730-00065)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BOARD OF PATENT APPEALS AND INTERFERENCES

Appellant(s): Ofir Zohar

Confirmation No.: 7108

Serial No.: 10/620,080

Filed: July 15, 2003

Title: **Data Allocation In a Distributed Storage System**

Examiner: Jasjit S. Vidwan

Group Art Unit: 2182

August 14, 2007

AMENDED APPEAL BRIEF FOR APPELLANTS

Board of Patent Appeals and Interferences
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

A Notice of Appeal was filed on February 28, 2007, along with a Pre-Appeal Request for Review. Appellant appeals to the Board of Patent Appeals and Interferences from the Office Action dated October 31, 2006, finally rejecting claims 1-10 and 23-32.

This communication is filed in response to the Notice of Non-Compliant Appeal Brief dated August 10, 2007. A response to the Notice is due September 10, 2007, and therefore, this communication is being timely filed.

Filed by Express Mail
Receipt No. EV479 107 17345
on 8-14-2007
pursuant to 37 C.F.R. 1.10.
By Frances Doyle
Frances Doyle

I. Real party in interest

The real party in interest is XIV Ltd., an Israeli corporation with offices at 1 Azrieli Center, Tel Aviv 67021, Israel.

II. Related appeals and interferences

Upon information and belief, there are no other appeals or interferences, which will directly affect, or be directly affected by, or have a bearing on the Board's decision in this appeal.

III. Status of claims

Claims 1-10 and 23-32 are pending, and claims 11-22 and 33-44 are withdrawn.

Claims 1-10 and 23-32 stand rejected and are appealed.

IV. Status of amendments

The final Office Action of October 31, 2006, does not object to any amendments previously made, and Appellant has not canceled or amended any of the claims in any filing since that Action.

V. Summary of claimed subject matter

Claim 1 relates to a method for data distribution that includes, *inter alia*, distributing logical addresses among an initial set of storage devices so as provide a balanced access to the devices (Specification, page 17, line 17, to page 18, line 10; Table 1; and Figure 2), and transferring the data to the storage devices in accordance with the logical addresses

(Specification; page 16, lines 24-30). The method of claim 1 further includes adding an additional storage device to the initial set, thus forming an extended set of the storage devices comprising the initial set and the additional storage device (Specification; page 23, lines 7-12). The method of claim 1 also includes redistributing the logical addresses among the storage devices in the extended set (Specification; page 23, lines 15-23) so as to cause a portion of the logical addresses to be transferred from the storage devices in the initial set to the additional storage device (Specification; page 23, lines 24-35; and Figure. 6), while maintaining the balanced access (Specification; page 25, lines 5-7), and while maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device (Specification; page 25, lines 7-11).

Claim 23 relates to a data distribution system that includes, *inter alia*, an initial set of storage devices among which are distributed logical addresses so as provide a balanced access to the devices (Specification, page 17, line 17, to page 18, line 10; Table 1; and Figure 2), and wherein data is stored in accordance with the logical addresses (Specification; page 16, lines 24-30). The system of claim 1 further includes an additional storage device to the initial set, thus forming an extended set of the storage devices comprising the initial set and the additional storage device (Specification; page 23, lines 7-12), the logical addresses being redistributed among the storage devices in the extended set (Specification; page 23, lines 15-23) so as to cause a portion of the logical addresses to be transferred from the storage devices in the initial set to the additional storage device (Specification; page 23, lines 24-35; and Figure. 6), while maintaining the balanced access (Specification; page 25, lines 5-7), and while maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device (Specification; page 25, lines 7-11).

VI. Grounds of rejection to be reviewed on appeal

1. Whether or not claims 1, 2, 5-10, 23, 24, and 27-32 are unpatentable under 35 U.S.C. 102(b) based on United States Patent No. 5,615,352 to Jacobson et al. (also referred to hereinafter as the “Jacobson” patent).

2. Whether or not claims 3, 4, 25, and 26 are unpatentable under 35 U.S.C. 103(a) based on the Jacobson patent in view of “Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the Worldwide Web,” by Karger et al., in the Proceedings of the 29th ACM Symposium on Theory of Computing, Pages 654-663 (also referred to hereinafter as the “Karger” reference).

VII. Argument

1. Claims 1, 2, 5-10, 23, 24, and 27-32 are not rendered unpatentable under 35 U.S.C. 102(b) based on the Jacobson patent

A. Claim 1 and its dependents are allowable over the Jacobson patent

Appellant respectfully traverses the rejection and asserts that the Jacobson patent does not disclose or suggest the feature of claim 1 of *maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device.*

Claim 1 relates to a method for data distribution that includes, *inter alia*, redistributing the logical addresses among the storage devices in the extended set so as to cause a portion of the logical addresses to be transferred from the storage devices in the initial set to the additional storage device. In claim 1, the redistributing is performed *while maintaining the balanced*

access and while maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device.

The Examiner asserts that the Jacobson patent discloses all of the features of claim 1.

However, the Jacobson patent does not disclose or suggest *maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device.*

The Examiner asserts that:

... the original logical addresses for the initial set of storage devices remain the same because Jacobson teaches that when more storage disks are added, *only the data that needs to be moved to the new disk is actually transferred* ... By moving only the data that needs to be transferred to additional storage device, the logical addresses for the data not moved would remain the same in the initial set of storage devices as it was prior to the addition of storage devices.

(Advisory Action; page 2, last paragraph; citing the Jacobson patent; col. 2, lines 9-25; emphasis added). However, this summary of the Jacobson patent, and in particular the high-lighted section, completely mischaracterizes the Jacobson patent. The cited section of the Jacobson patent actually states only that:

According to one method, the physical storage space of the disk array is configured into multiple stripes for storing a predetermined amount of data. The stripes extend across multiple storage disks in the disk array, and are made up of one or more equal sized segments from each storage disk in the disk array. *When more storage disks are added, data from one stripe is moved to a another portion of the physical storage space. The physical storage space containing the stripe is then reconfigured into an expanded stripe for storing data.* This expanded stripe spans across all storage disks, including the new additional storage disks. The expanded stripe is then ready to receive new data. This process is continued stripe-by-stripe until all stripes have been configured to include the new disks. While the stripe is being expanded, the data storage system does not allocate any virtual blocks within the selected stripe.

(The Jacobson patent; col. 2, lines 9-25; emphasis added). As is readily apparent from the cited section, the Jacobson patent does NOT teach that only data that needs to be moved to a new disk is actually transferred. The Jacobson patent indicates that data is moved when storage disks are added to the system, and there is no suggestion that the “another portion of the physical disk storage space” discussed in the section of the Jacobson patent cited above is limited in any way to the new storage disk. Therefore, there is no disclosure or suggestion in the Jacobson patent of maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device, as recited in claim 1.

More importantly, a review of Figures 7 and 8 of the Jacobson patent confirms unequivocally that the Jacobson patent discloses, upon adding additional storage, moving logical addresses to storage devices that were part of the original storage array as part of a reconfiguring. Figure 7 illustrates reorganizing data upon addition of a new storage device, disk 4, which creates new space for every group. As is apparent from figure 7, data from segment 3 in Group B, which is initially in disk 0, is **moved** to Group A, in disk 3. Disk 3 is NOT a new storage device, but rather is part of the *original* storage array. Therefore, upon adding new disk 4, data is **moved from an initial position to a new position which is NOT in the newly added space**, but rather is in the initial space. This arrangement is described in detail in the Jacobson patent, which states that:

[n]o reorganization is necessary. Instead, data **from segment 3 of group B is simply moved** to fill segment P in group A that **had previously held parity information.**

(The Jacobson patent; col. 14, lines 38-41; emphasis added; see also col. 14, lines 54-62).

Similarly, Figure 8 of the Jacobson patent illustrates four data segment that are moving (4, 5, 6, and 7) and **only one** of them, data segment 7, is directed to the new space. All of the other three

data segments, segments 4, 5, and 6, migrate from an initial space in the storage array to *a different space in the original storage array.*

The Jacobson patent apparently adds disk space and redistributes the addresses (though not necessarily in a load-balanced manner) so that data from one stripe may be moved to another portion of the original physical storage space. There is no limitation presented in the Jacobson patent that data that is in a storage area prior to expansion is *not moved* to another storage area which was also part of the original storage area. In contrast, Applicant's invention recites moving data *only* into the added space, and never moving data within the original storage space after expanding the storage area. The Jacobson patent moves data within the storage space, but not necessarily *only* to the added space. Therefore, the Jacobson patent does not identically disclose all of the features of claim 1, and therefore claim 1 is allowable over the reference.

Additionally, the Jacobson patent does not disclose or suggest maintaining a balanced access to the devices after the addition of a storage device. There is no disclosure in any of the cited sections of the Jacobson patent of performing *load balancing* and *redistributing*. Therefore, for at least this additional reason claim 1 is allowable over the Jacobson patent.

Claims 2 and 5-10 depend from claim 1 and therefore include all of the features discussed above in support of the patentability of claim 1. Therefore these claims are allowable at least for the same reasons that claim 1 is allowable.

B. Claim 23 and its dependents are allowable over the Jacobson patent

Claim 23 relates to a data distribution system that includes, *inter alia*, an initial set of storage devices among which are distributed logical addresses so as provide a balanced access to the devices, and wherein data is stored in accordance with the logical addresses. The data

distribution system of claim 23 also includes an additional storage device to the initial set, thus forming an extended set of the storage devices comprising the initial set and the additional storage device, the logical addresses being redistributed among the storage devices in the extended set so as to cause a portion of the logical addresses to be transferred from the storage devices in the initial set to the additional storage device, *while maintaining the balanced access and while maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device.*

As discussed above, the Jacobson patent redistributes at least some data stored in the original set of storage devices to other devices within the original set of storage devices after the addition of a new storage device. Therefore, there is no disclosure or suggestion in the Jacobson patent of *maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device*, as recited in claim 23. Additionally, the Jacobson patent does not disclose or suggest maintaining a balanced access to the devices after the addition of a storage device. There is no disclosure in any of the cited sections of the Jacobson patent of performing *load balancing* and *redistributing*. Therefore, for at least this additional reason claim 23 is allowable over the Jacobson patent.

Claims 24 and 27-32 depend from claim 23 and therefore include all of the features discussed above in support of the patentability of claim 23. Therefore these claims are allowable at least for the same reasons that claim 23 is allowable.

2. Claims 3, 4, 25, and 26 are not rendered unpatentable under 35 U.S.C. 103(a) based on the Jacobson patent in view of the Karger reference


The Karger reference fails to disclose or suggest the features of independent claims 1 and 23 of redistributing data after the addition of a storage disk *while maintaining the balanced*

access and while maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device. Claims 3 and 4 depend from claim 1 and claims 25 and 26 depend from claim 23, and therefore each of these claims includes all of the features of the claim from which it depends. Therefore each of these claims is allowable for at least the same reasons as their respective base claim is allowable.

CONCLUSION

Claims 1, 2, 5-10, 23, 24, and 27-32 are patentable over the Jacobson patent, and claims 3, 4, 25, and 26 are patentable over the Jacobson patent in view of the Karger reference. Accordingly, it is respectfully submitted that the Examiner erred in rejecting claims 1-10 and 23-32, and a reversal of such rejections by this Honorable Board is solicited.

Respectfully submitted,



Brian E. Hennessey
Reg. No. 51,271

CUSTOMER NUMBER 026304
Telephone: (212) 940-6311
Fax: (212) 940-8986/8987
Docket No.: ASSIA 20.502 (056730-00065)
BEH:fd

VIII. Claims Appendix

1. A method for data distribution, comprising:
 - distributing logical addresses among an initial set of storage devices so as provide a balanced access to the devices;
 - transferring the data to the storage devices in accordance with the logical addresses;
 - adding an additional storage device to the initial set, thus forming an extended set of the storage devices comprising the initial set and the additional storage device; and
 - redistributing the logical addresses among the storage devices in the extended set so as to cause a portion of the logical addresses to be transferred from the storage devices in the initial set to the additional storage device, while maintaining the balanced access and while maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device.
2. A method according to claim 1, wherein redistributing the logical addresses comprises no transfer of the logical addresses between the storage devices in the initial set.
3. A method according to claim 1, wherein distributing the logical addresses comprises applying a consistent hashing function to the initial set of storage devices so as to determine respective initial locations of the logical addresses among the initial set, and wherein redistributing the logical addresses comprises applying the consistent hashing function to the extended set of storage devices so as to determine respective subsequent locations of the logical addresses among the extended set.

4. A method according to claim 1, wherein distributing the logical addresses comprises applying a randomizing function to the initial set of storage devices so as to determine respective initial locations of the logical addresses among the initial set, and wherein redistributing the logical addresses comprises applying the randomizing function to the extended set of storage devices so as to determine respective subsequent locations of the logical addresses among the extended set.
5. A method according to claim 1, wherein at least one of the storage devices comprises a fast access time memory.
6. A method according to claim 1, wherein at least one of the storage devices comprises a slow access time mass storage device.
7. A method according to claim 1, wherein the storage devices have substantially equal capacities, and wherein distributing the logical addresses comprises distributing the logical addresses substantially evenly among the initial set, and wherein redistributing the logical addresses comprises redistributing the logical addresses substantially evenly among the extended set.
8. A method according to claim 1, wherein a first storage device comprised in the storage devices has a first capacity different from a second capacity of a second storage device comprised in the storage devices, and wherein distributing the logical addresses comprises distributing the logical addresses substantially according to a ratio of the first capacity to the

second capacity, and wherein redistributing the logical addresses comprises redistributing the logical addresses substantially according to the ratio.

9. A method according to claim 1, wherein distributing the logical addresses comprises allocating a specific logical address to a first storage device and to a second storage device, the first and second storage devices comprising different storage devices, and wherein storing the data comprises storing a first copy of the data on the first storage device and a second copy of the data on the second storage device.

10. A method according to claim 1, and comprising writing the data from a host external to the storage devices, and reading the data to the external host from the storage devices.

23. A data distribution system, comprising:

an initial set of storage devices among which are distributed logical addresses so as provide a balanced access to the devices, and wherein data is stored in accordance with the logical addresses; and

an additional storage device to the initial set, thus forming an extended set of the storage devices comprising the initial set and the additional storage device, the logical addresses being redistributed among the storage devices in the extended set so as to cause a portion of the logical addresses to be transferred from the storage devices in the initial set to the additional storage device, while maintaining the balanced access and while maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device.

24. A system according to claim 23, and wherein the logical addresses are redistributed so that there is no transfer of the logical addresses between the storage devices in the initial set.
25. A system according to claim 23, wherein the distributed logical addresses are determined by applying a consistent hashing function to the initial set of storage devices so as to determine respective initial locations of the logical addresses among the initial set, and wherein redistributing the logical addresses comprises applying the consistent hashing function to the extended set of storage devices so as to determine respective subsequent locations of the logical addresses among the extended set.
26. A system according to claim 23, wherein the distributed logical addresses are determined by applying a randomizing function to the initial set of storage devices so as to determine respective initial locations of the logical addresses among the initial set, and wherein redistributing the logical addresses comprises applying the randomizing function to the extended set of storage devices so as to determine respective subsequent locations of the logical addresses among the extended set.
27. A system according to claim 23, wherein at least one of the storage devices comprises a fast access time memory.
28. A system according to claim 23, wherein at least one of the storage devices comprises a slow access time mass storage device.

29. A system according to claim 23, wherein the storage devices have substantially equal capacities, and wherein the distributed logical addresses are distributed substantially evenly among the initial set, and wherein redistributing the logical addresses comprises redistributing the logical addresses substantially evenly among the extended set.

30. A system according to claim 23, wherein a first storage device comprised in the storage devices has a first capacity different from a second capacity of a second storage device comprised in the storage devices, and wherein the distributed logical addresses are distributed substantially according to a ratio of the first capacity to the second capacity, and wherein redistributing the logical addresses comprises redistributing the logical addresses substantially according to the ratio.

31. A system according to claim 23, wherein the distributed logical addresses comprise a specific logical address allocated to a first storage device and to a second storage device, the first and second storage devices comprising different storage devices, and wherein storing the data comprises storing a first copy of the data on the first storage device and a second copy of the data on the second storage device.

32. A system according to claim 23, and comprising a memory having a table wherein is stored a correspondence between a plurality of the logical addresses and a specific storage device in the initial set, wherein the plurality of the logical addresses are related to each other by a mathematical relation.

IX. Evidence Appendix

No evidence was submitted to or entered by the Examiner during prosecution of this application.

X. Related Proceedings Appendix

No appeals or interferences, which will directly affect or be directly affected by or have a bearing on the Board's decision in this appeal, are identified in section II of this brief, and therefore there are no decisions rendered by a court or the Board in any proceeding included in this section.